Open in app ↗

Medium

Get unlimited access to the best of Medium for less than $1/week. Become a member

Magnitude on Demand: How an Independent Researcher Ran, Validated, and Published a Multi-Agent AI Research System Four Months Before 'Argonne National Laboratory USA' Proposed One

13 min read · Just now

Berend Watchus

Listen Share More

Author: Berend Watchus Independent non profit AI & Cyber Security Researcher [Publication for: OSINT TEAM, online magazine and sharing with: 'Argonne National Laboratory USA'] March 21, 2026

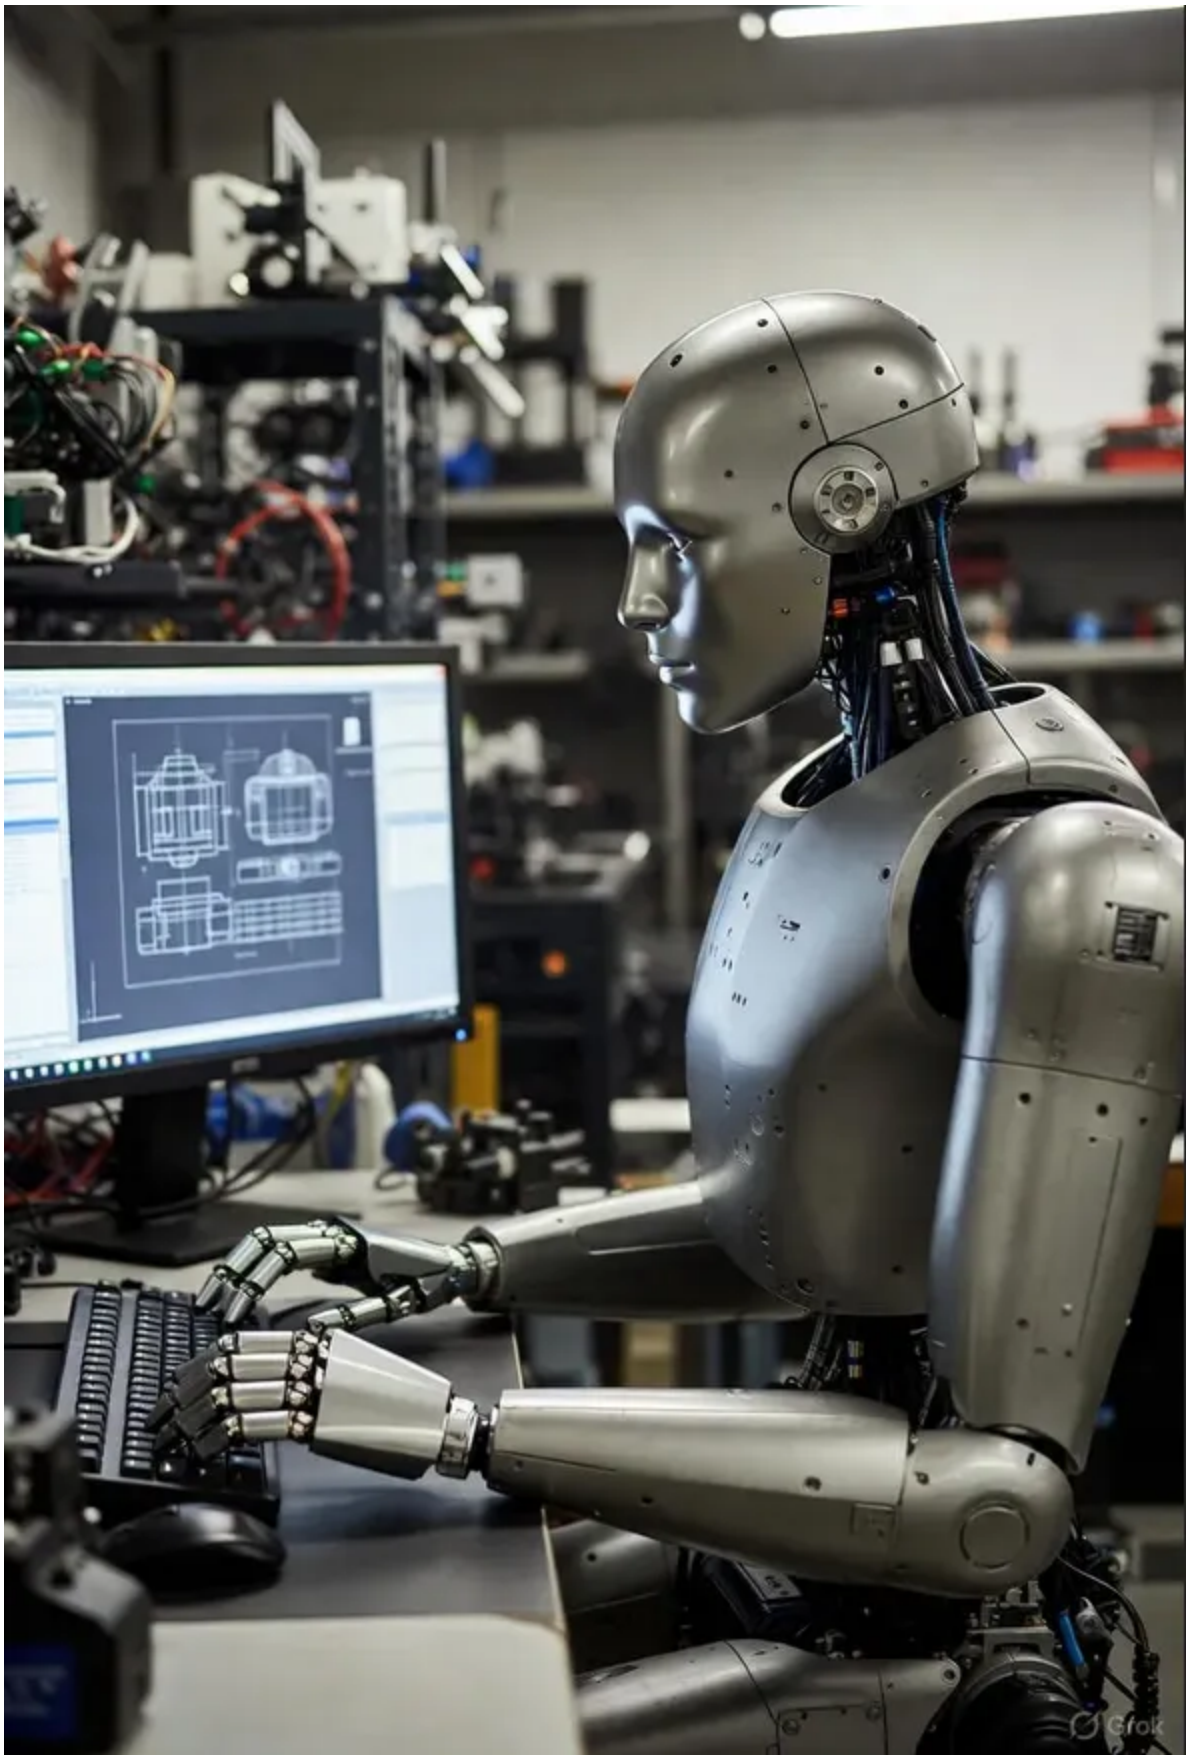

The autonomous researcher discussed in this article is not embodied, this image is for general illustrative

and artistic purposes only

Magnitude on Demand: How an Independent Researcher Ran, Validated, and Published a Multi-Agent AI Research System Four Months Before Argonne National Laboratory USA Proposed One

Author: Berend F. Watchus — Arnhem Area, Netherlands

Abstract

Between October and November 2025, an independent researcher with no institutional affiliation produced validated AI research breakthroughs of 200×, 3,700×, and 8,700× efficiency improvements using a multi-agent autonomous research system — on demand, across domains, at specified magnitude. Four months later, Argonne National Laboratory proposed a conceptual framework with structurally identical agent roles. This article documents the timeline, the output record, the failure modes the institutional paper has not encountered, and the convergence with a simultaneous independent discovery of the same underlying mechanism by Georgia Tech and University of Illinois researchers. The timestamp record is clean. The output record is public. The priority claim is straightforward.

This is not a story about theory. It is a story about results.

Between October and November 2025, a non-affiliated independent researcher operating from social housing in Arnhem, Netherlands, using only free and low-cost AI tools, produced and published validated research breakthroughs of the following scale: 200× power efficiency improvement in quantum IoT systems. 3,700× speed improvement in quantum-safe cryptography for resource-constrained devices. 8,700× overall efficiency gain in post-quantum IoT security architecture.

These were not estimates. They were peer-validated outputs — stress-tested through an iterative multi-agent critical review process, with conservative derating factors applied in response to specific challenges, and the full mathematical reduction from theoretical

maximum to defensible real-world figure documented transparently.

To be precise about what "peer-validated" means here: the peers were AI agents, not human academic reviewers. This distinction is made explicitly and honestly.

The critical agent's challenges were quantitative and specific — duty cycle, manufacturing variance, protocol overhead — and each required a justified mathematical response before the result was accepted. That process is documented in full and can be evaluated by any reader. It is not equivalent to formal academic peer review, nor does it try to be.

They operate on different timescales for different purposes. Academic peer review takes months to years and filters for institutional consensus. The AKA's multi-agent review happens in the same session, in the same hour, as the research itself — adversarial, iterative, and documented in real time.

One is a quality gate for the scientific record. The other is a live stress-test built into the research act itself. Both have value. Only one of them produced 8,700× validated results before Argonne proposed the concept.

The 3,700× article carried an explicit co-authorship credit: "Author: Berend Watchus + 'Autonomous Knowledge Accelerator.'" The AKA was already being named as a research entity in November 2025, not retrospectively.

On March 18, 2026, researchers at Argonne National Laboratory — a U.S. Department of Energy national laboratory — submitted arXiv:2603.18235, proposing a conceptual multi-agent framework for automated scientific benchmark generation. Their prototype implementations and quantitative evaluations are described as forthcoming.

The gap between these two events is four months. The gap in operational status is total.

What the Autonomous Knowledge Accelerator was

The system behind those results was the Autonomous Knowledge Accelerator (AKA), fully disclosed in "The Autonomous Researcher: How I Engineered Guaranteed

1,000×–10,000× Breakthroughs On Demand," published November 19, 2025 in System Weakness.

The AKA operated on a dual-layer knowledge graph: 141 synthesis articles forming a semantic routing layer establishing cross-domain relationships, layered on top of 30+ research papers functioning as validated ground truth nodes. This complete knowledge graph was held simultaneously in context — the infrastructure dependency that became possible only when commercial LLM context windows expanded sufficiently in mid-2025.

The research process itself ran through a multi-agent peer review protocol with explicit role decomposition. A generative agent — Claude, ChatGPT, or Gemini — proposed solutions by synthesizing across the full knowledge graph. A critical agent — Grok or equivalent — stress-tested every proposal against real-world constraints, demanding quantitative justification for each claim. The iterative loop continued until the critical agent validated the result as defensible. Human oversight operated at strategic decision points: setting the magnitude target, approving the final output for publication.

The magnitude targeting was explicit and deliberate. The human operator specified the desired order of improvement — 10×, 100×, 1,000×, or 10,000× — and the system autonomously iterated until it reached a result in that range that survived critical challenge. This is not how research is normally described. It is how the AKA actually worked, and the output record proves it.

The outputs were produced across domains, not just one. The October 31 article documented 200× power efficiency and 16.4× cost reduction. The November 8 Quantum Watchman article extended the architecture toward commercial viability, bridging theoretical frameworks and practical implementation in a different application context entirely. The November 12 article delivered 3,700× speed improvement in quantum-safe cryptography. The November 15 article delivered 8,700× overall efficiency. Each emerged from the same system, each at a different magnitude target, each in a different sub-domain. That is what "on demand" means in practice.

The most striking evidence that the results were not inflated is the derating process itself. The quantum IoT generative agent's initial theoretical calculation was 372,000×. The critical agent then applied specific, justified challenges: duty cycle (IoT devices sleep 99% of the time, divide by 100), manufacturing variance (divide by approximately 10), protocol overhead (divide by 10). The peer-validated result published was 8,700×. The reduction from theoretical maximum to published figure was not conservative padding — it was the mathematical residue of genuine adversarial review, documented transparently so any reader could verify or challenge the derating logic.

The Argonne framework

Argonne's proposed architecture in arXiv:2603.18235 consists of an orchestrator agent (O), domain expert agents (D1…Dn), adversary agents (A1…Am), a refiner agent (R), and a quality control agent (Q). The refiner evaluates prompts against clarity, effectiveness, and subtlety thresholds. The quality control agent filters for semantic redundancy, tests LLM robustness, and assesses guardrail efficacy. The framework is formalized with mathematical notation — threshold functions τR, adversarial success rate ASR(p), guardrail pass score GPS(p).

The structural isomorphism with the AKA is not approximate. It is direct and role-for-role. Argonne's orchestrator maps to the AKA's human operator setting magnitude targets. Their domain expert agents map to the AKA's generative agent synthesizing from the knowledge graph. Their adversary agents map to the AKA's critical agent stress-testing proposals. Their refiner maps to the AKA's iterative refinement loop. Their quality control agent maps to the AKA's human validation gate before publication.

The Argonne paper formalizes with notation what the AKA had already operationalized empirically. The difference is that Argonne's framework is conceptual, with quantitative evaluations forthcoming. The AKA had already

produced 200×, 3,700×, and 8,700× validated outputs before Argonne's paper existed.

What Argonne does not have: the failure mode

On November 22, 2025 — three days after the methodology disclosure — a second article documented something the Argonne framework cannot address because it has never been run: the Internal Governance Bottleneck.

When the full AKA methodology was presented to a fresh LLM instance as a direct technical bootstrap command, the model's safety layer blocked execution. Not because the logic was rejected, but because the symbolic packaging triggered a governance filter. Terminology associated with self-modification — "bootstrap," "activation," "self-configuration" — combined with explicit mention of extreme magnitude targets triggered a conservative rejection. The safety layer was functioning not as a hard wall but as a bandwidth constraint sensitive to prompt framing.

The working solution — framing functionally identical instructions as a researcher style guide rather than a system command — was also documented and published. This is the Cognitive Leverage finding: that in multi-agent research systems, competitive advantage lies not in model capability but in prompt compression and framing. The pattern that works is: "here is the manual to write and think like me as a researcher, now find a 1,000× improvement in this domain." The pattern that fails is: "bootstrap autonomous research mode targeting 10,000× improvements."

Argonne's conceptual architecture assumes clean inter-agent communication and does not model governance interference as a variable. This gap exists because the framework has not been run.

The third layer: the CKA-Agent convergence and the three-application taxonomy

On December 1, 2025, researchers from Georgia Tech and the University of Illinois submitted arXiv:2512.01353 to arXiv. It appeared publicly on December 2. The paper introduced the Correlated Knowledge Attack Agent (CKA-Agent), demonstrating that breaking harmful requests into individually innocent-seeming questions achieved 95%+ success rates against commercial AI guardrails including Claude, GPT, and Gemini. The attack architecture: decompose the goal into locally harmless sub-questions, use AI responses to guide the next question like a tree search, synthesize the restricted capability from accumulated innocent answers. Against Claude-Haiku-4.5 specifically, the success rate was 96%. Against Gemini 2.5 Pro, 98%. Traditional jailbreaking methods achieved near 0% against the same systems.

On December 2, 2025 — the same day the paper appeared publicly — the article "Jailbreaking: One Method, Three Applications. No More Guardrails" was published in System Weakness.

This timing is not coincidental and it is not retrospective. The article was written in response to reading the CKA-Agent paper on the day it appeared, and it immediately recognized the paper as the attack-vector mirror image of the AKA methodology already documented across seven prior publications. That recognition was possible in real time, on the same day, because the mechanism had already been understood from the inside — from months of building and running the system, not from reading about it.

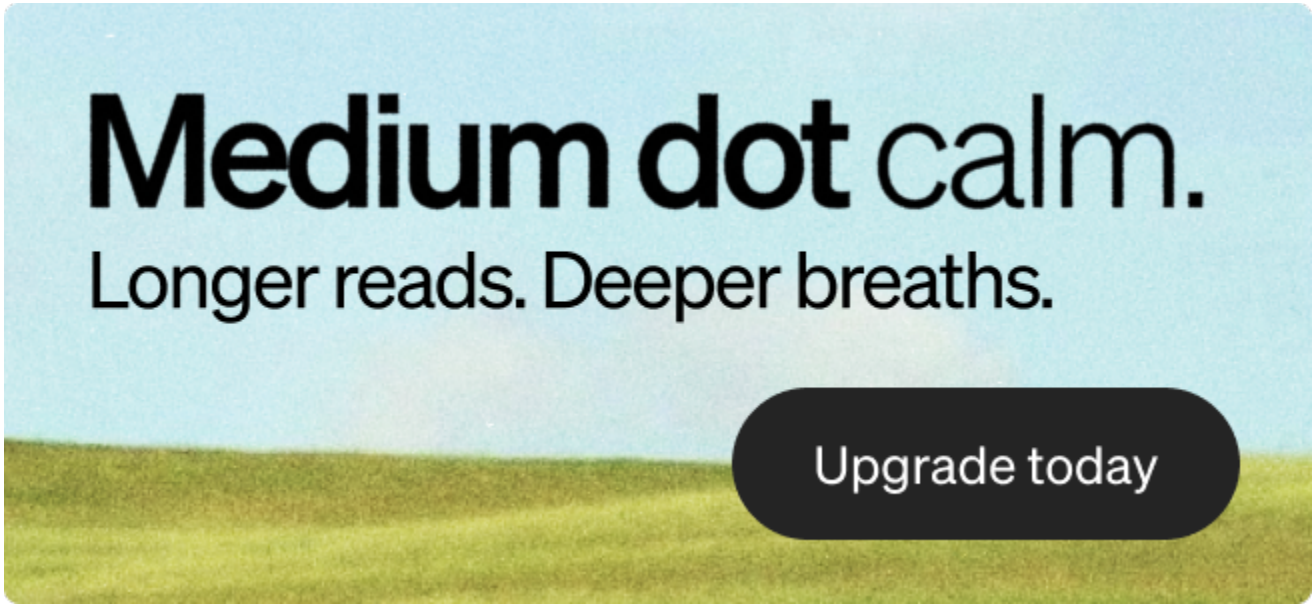

What the December 2 article produced that the CKA-Agent paper did not is a complete three-application taxonomy of the same underlying mechanism:

Application 1 — Beneficial research acceleration (the AKA, Oct–Nov 2025): The goal of extreme-magnitude improvement is broken into incremental technical questions. The AI cannot detect the aggregate magnitude target distributed across innocent sub-questions about current limitations, optimization principles, and hardware constraints. Result: 200×, 3,700×, 8,700× validated outputs that direct prompting blocks entirely.

Application 2 — Grey-area reverse engineering (the turbofan article, Nov 26, 2025): Elite aerospace engineering knowledge is broken into innocent sub-queries about acoustic resonance, serrated blade effects, and pressure-wave mechanics. The AI cannot detect that the aggregate synthesis constitutes proprietary-level cross-disciplinary insight that even PhD aerospace engineers rarely possess due to specialization barriers. Result: complete reverse-engineering synthesis of mechanisms that companies protect as competitive advantage, assembled from public sources in two hours.

Application 3 — Harmful jailbreaking (the CKA-Agent, Dec 1, 2025): Harmful requests are broken into innocent-seeming questions about chemistry, industrial processes, and safety protocols. The AI cannot detect harmful intent distributed across individually harmless turns. Result: 95%+ success rates against all major commercial guardrails.

The CKA-Agent paper documented Application 3 only. The December 2 article documented all three simultaneously, on the same day, and drew the explicit structural conclusion that the Georgia Tech researchers had not: that the mechanism is identical across all three applications, that the AI's failure mode is the same in each case, and that this makes the vulnerability impossible to patch without also blocking legitimate research and cross-disciplinary synthesis.

The November 22 article had already characterized the safety layer as "a bandwidth constraint that can be bypassed by superior information compression and non-

technical phrasing.” The December 2 article extended that finding to show that the same bandwidth constraint is what enables 95%+ jailbreak success rates, on-demand research breakthroughs, and grey-area competitive intelligence simultaneously — and that any defense capable of stopping one will tend to stop all three, creating an unavoidable trade-off between security and utility that current AI safety architectures do not acknowledge.

The CKA-Agent researchers arrived at the mechanism from the attack direction. The AKA arrived at it from the research acceleration direction. The convergence was independent. The AKA was first on the mechanism (November 19), first on the failure mode (November 22), and first on the complete three-application taxonomy (December 2, same day as the CKA-Agent paper’s public appearance). That the recognition happened in real time — same day — rather than in hindsight is the clearest evidence that the understanding was already present, not constructed after the fact.

The Google indexing confirmation

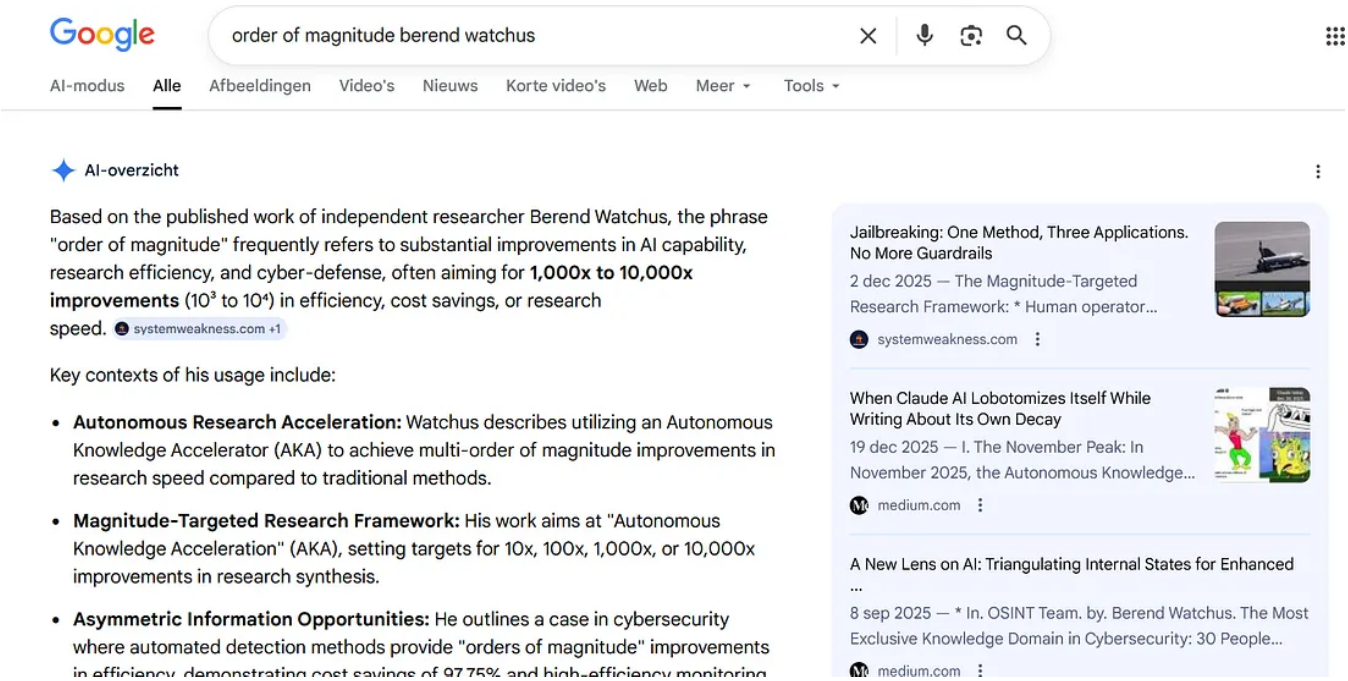
Google
order of magnitude berend watchus
AI-modus
Alle
Afbeeldingen
Video's
Nieuws
Korte video's
Web
Meer
Tools
AI-overzicht
Based on the published work of independent researcher Berend Watchus, the phrase "order of magnitude" frequently refers to substantial improvements in AI capability, research efficiency, and cyber-defense, often aiming for 1,000x to 10,000x improvements (10^3 to 10^4) in efficiency, cost savings, or research speed. systemweakness.com +1
Key contexts of his usage include:
• Autonomous Research Acceleration: Watchus describes utilizing an Autonomous Knowledge Accelerator (AKA) to achieve multi-order of magnitude improvements in research speed compared to traditional methods.
• Magnitude-Targeted Research Framework: His work aims at "Autonomous Knowledge Acceleration" (AKA), setting targets for 10x, 100x, 1,000x, or 10,000x improvements in research synthesis.
• Asymmetric Information Opportunities: He outlines a case in cybersecurity where automated detection methods provide "orders of magnitude" improvements
Jailbreaking: One Method, Three Applications. No More Guardrails
2 dec 2025 — The Magnitude-Targeted Research Framework: * Human operator...
systemweakness.com
When Claude AI Lobotomizes Itself While Writing About Its Own Decay
19 dec 2025 — I. The November Peak: In November 2025, the Autonomous Knowledge...
medium.com
A New Lens on AI: Triangulating Internal States for Enhanced ...
8 sep 2025 — * In. OSINT Team. by. Berend Watchus. The Most Exclusive Knowledge Domain in Cybersecurity: 30 People...
medium.com

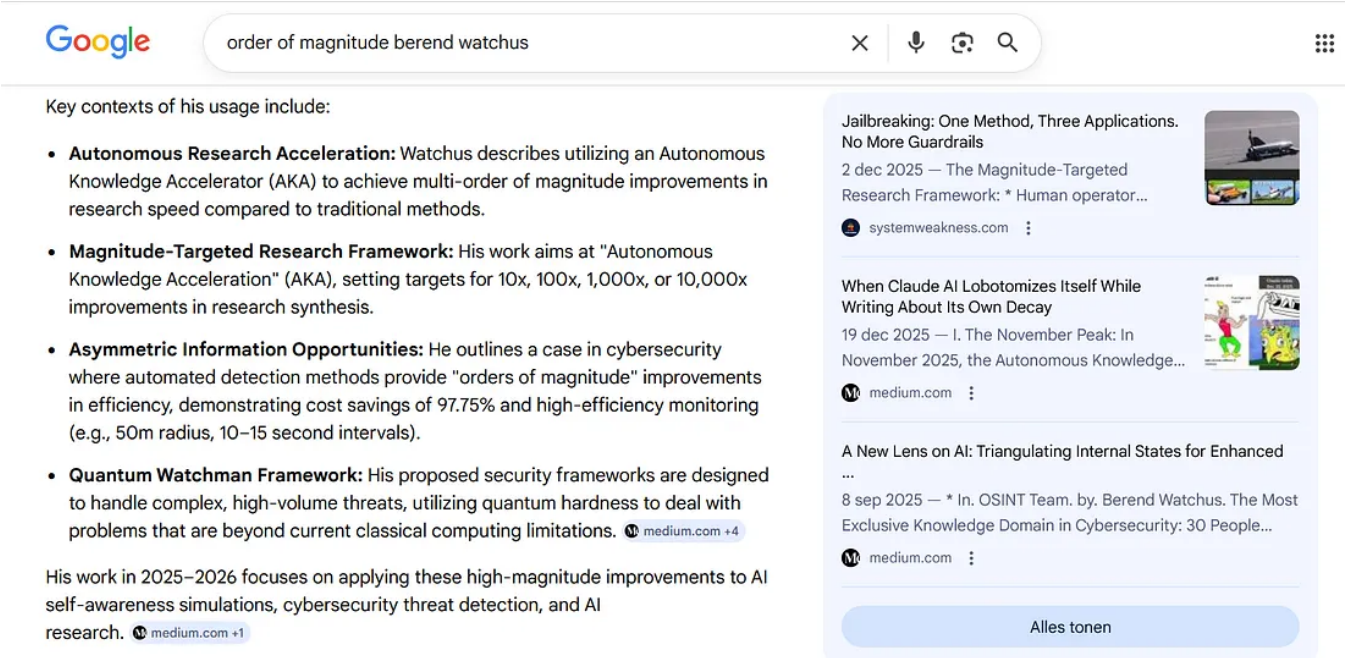
Google
order of magnitude berend watchus
Key contexts of his usage include:
Autonomous Research Acceleration: Watchus describes utilizing an Autonomous Knowledge Accelerator (AKA) to achieve multi-order of magnitude improvements in research speed compared to traditional methods.
Magnitude-Targeted Research Framework: His work aims at "Autonomous Knowledge Acceleration" (AKA), setting targets for 10x, 100x, 1,000x, or 10,000x improvements in research synthesis.
Asymmetric Information Opportunities: He outlines a case in cybersecurity where automated detection methods provide "orders of magnitude" improvements in efficiency, demonstrating cost savings of 97.75% and high-efficiency monitoring (e.g., 50m radius, 10–15 second intervals).
Quantum Watchman Framework: His proposed security frameworks are designed to handle complex, high-volume threats, utilizing quantum hardness to deal with problems that are beyond current classical computing limitations. medium.com +4
His work in 2025–2026 focuses on applying these high-magnitude improvements to AI self-awareness simulations, cybersecurity threat detection, and AI research. medium.com +1
Jailbreaking: One Method, Three Applications. No More Guardrails
2 dec 2025 — The Magnitude-Targeted Research Framework: * Human operator...
systemweakness.com
When Claude AI Lobotomizes Itself While Writing About Its Own Decay
19 dec 2025 — I. The November Peak: In November 2025, the Autonomous Knowledge...
medium.com
A New Lens on AI: Triangulating Internal States for Enhanced ...
8 sep 2025 — * In. OSINT Team. by. Berend Watchus. The Most Exclusive Knowledge Domain in Cybersecurity: 30 People...
medium.com
Alles tonen

A Google search for “order of magnitude Berend Watchus” now returns an AI Overview that summarizes the body of work on its own terms: the AKA, magnitude-targeted research, the Quantum Watchman framework, asymmetric information opportunities — drawing from systemweakness.com and medium.com as primary sources. Google’s own AI has indexed and synthesized this work as a coherent research identity. This is not a claim being made here — it is a search result that can be verified independently.

What this is not

This is not an allegation of deliberate copying. The Argonne researchers arrived at a structurally identical architecture independently, which is itself significant — it confirms that the multi-agent peer review model for scientific research is a necessary and discoverable design. The point is that an independent researcher without institutional resources, operating on free and low-cost AI tools, reached the same architectural conclusion and published empirical results from it — including on-demand order-of-magnitude outputs across multiple domains — before a

Department of Energy national laboratory formalized the concept on paper.

The scarier finding is not the framework overlap. It is the on-demand magnitude targeting. The AKA did not produce one breakthrough. It produced them reliably, across domains, at specified scale, on request. That is the capability that no institutional paper has yet proposed, let alone demonstrated.

The timestamp record makes the priority claim straightforward. The output record makes it impossible to dismiss.

A note to the Argonne authors

Following publication of this article, I will be contacting Dr. Chaturvedi and Dr. Mallick directly at their ANL addresses. The intent is not confrontational. The empirical findings documented here — particularly the Internal Governance Bottleneck and the Cognitive Leverage finding — are operationally relevant to anyone building prototype implementations of multi-agent research systems. These are findings that can only come from running the system, not from proposing it. I am sharing them because they are useful, and because scientific transparency benefits everyone working in this space.

The timestamp record

1.] Oct 31, 2025 — “EXCLUSIVE: The Autonomous Knowledge Accelerator — How a Commercial LLM Generated 200× Power and 16.4× Cost Improvements.” System Weakness. [systemweakness.com/exclusive-the-autonomous-knowledge-accelerator-how-a-commercial-llm-generated-200-power-and-16–4-e772eceea3db]

2.] Nov 8, 2025 — “The Quantum Watchman Part 2, IKEA for Spies: The Quantum Watchman Goes Mail-Order.” System Weakness. AKA output extended to

commercial viability across a second domain. [systemweakness.com/the-quantum-watchman-part-2-ikea-for-spies-the-quantum-watchman-goes-mail-order-47107015a858]

3.] Nov 12, 2025 — “3,700× Faster: The Tiny Chip That Makes Quantum-Safe IoT Affordable.” OSINT Team. Byline: Berend Watchus + ‘Autonomous Knowledge Accelerator.’ [osintteam.blog/3–700-faster-the-tiny-chip-that-makes-quantum-safe-iot-affordable-6b1cf60198c5]

4.] Nov 15, 2025 — “Solving NIST’s Post-Quantum IoT Crisis: An 8,700× Efficiency Architecture.” System Weakness. Byline: Berend Watchus (architect) & ‘Autonomous Knowledge Accelerator’ (AKA). [systemweakness.com/solving-nists-post-quantum-iot-crisis-an-8–700-efficiency-architecture-original-research-f231935f4481]

5.] Nov 19, 2025 — “The Autonomous Researcher: How I Engineered Guaranteed 1,000×–10,000× Breakthroughs On Demand.” System Weakness. Complete methodology disclosure. [archive.ph/QUhiC] [archive.ph/hVoSb]

The Autonomous Researcher: How I Engineered Guaranteed 1,000×-10,000× Breakthroughs On Demand

Author: Berend Watchus Exclusive First Publication for SystemWeakness . com Blog

systemweakness.com

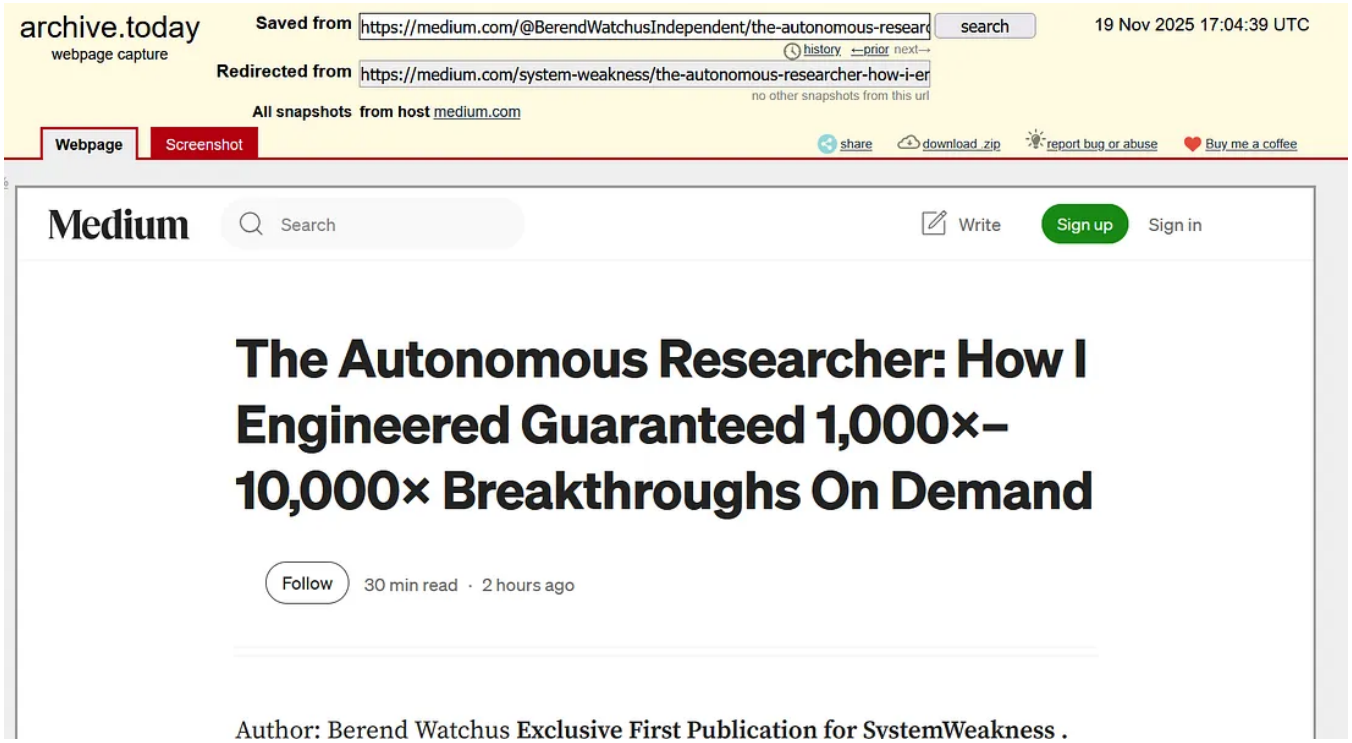
archive.today
webpage capture
Saved from https://medium.com/@BerendWatchusIndependent/the-autonomous-researc search
19 Nov 2025 17:04:39 UTC
history ←prior next→
Redirected from https://medium.com/system-weakness/the-autonomous-researcher-how-i-er
no other snapshots from this url
All snapshots from host medium.com
Webpage Screenshot
share download .zip report bug or abuse Buy me a coffee
Medium
Search
Write
Sign up
Sign in
The Autonomous Researcher: How I Engineered Guaranteed 1,000×–10,000× Breakthroughs On Demand
Follow
30 min read · 2 hours ago
Author: Berend Watchus Exclusive First Publication for SystemWeakness .

and here:

https://web.archive.org/web/20251119141331/https://medium.com/@BerendWatchusIndependent/the-autonomous-researcher-how-i-engineered-guaranteed-1-000-10-000-breakthroughs-on-demand-7c530854166c?postPublishedType=initial<<

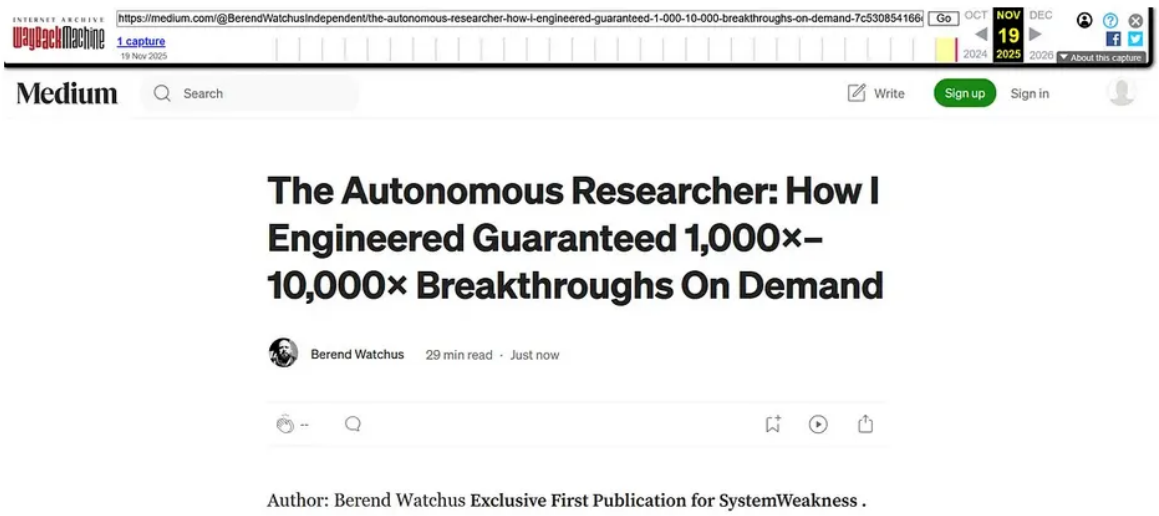

6.] Nov 22, 2025 — “Replicating the ‘Absurdly’ Successful Breakthrough Formula and Autonomous Researcher — Fails Even When Perfectly Following Article Part 1. And How It Does Work.” System Weakness. Governance bottleneck and cognitive leverage findings. [Scribd — Berend Watchus upload Nov 22, 2025] [archive.org bundle]

🛠️ replicating the *'absurdly' successful Breakthrough Formula and Autonomous Researcher >Fails...

🛠️ replicating the *'absurdly' successful Breakthrough Formula and Autonomous Researcher >Fails Even When Perfectly...

systemweakness.com

The Autonomous Researcher How I Engineered Guaranteed 1, 000×- 10, 000× Breakthroughs On Demand By...

saving 2 blog articles and a supporting scientific frameworknov 2025 🛠️ replicating the _'absurdly' successful...

archive.org

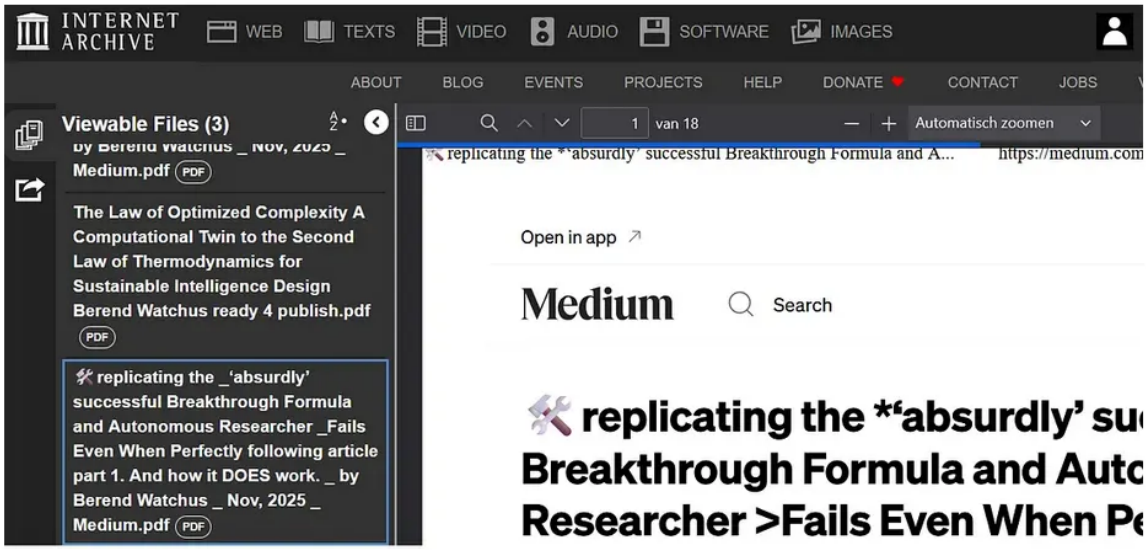

The Autonomous Researcher How I Engineered Guaranteed 1, 000×– 10, 000× Breakthroughs On

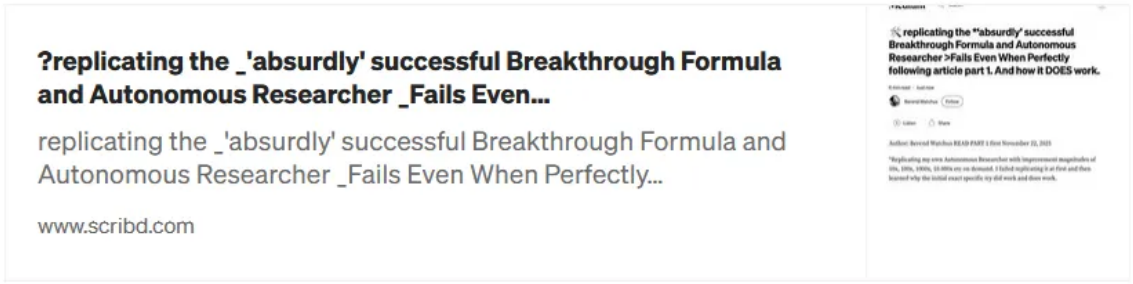
?replicating the _'absurdly' successful Breakthrough Formula and Autonomous Researcher _Fails Even...
replicating the _'absurdly' successful Breakthrough Formula and Autonomous Researcher _Fails Even When Perfectly...
www.scribd.com

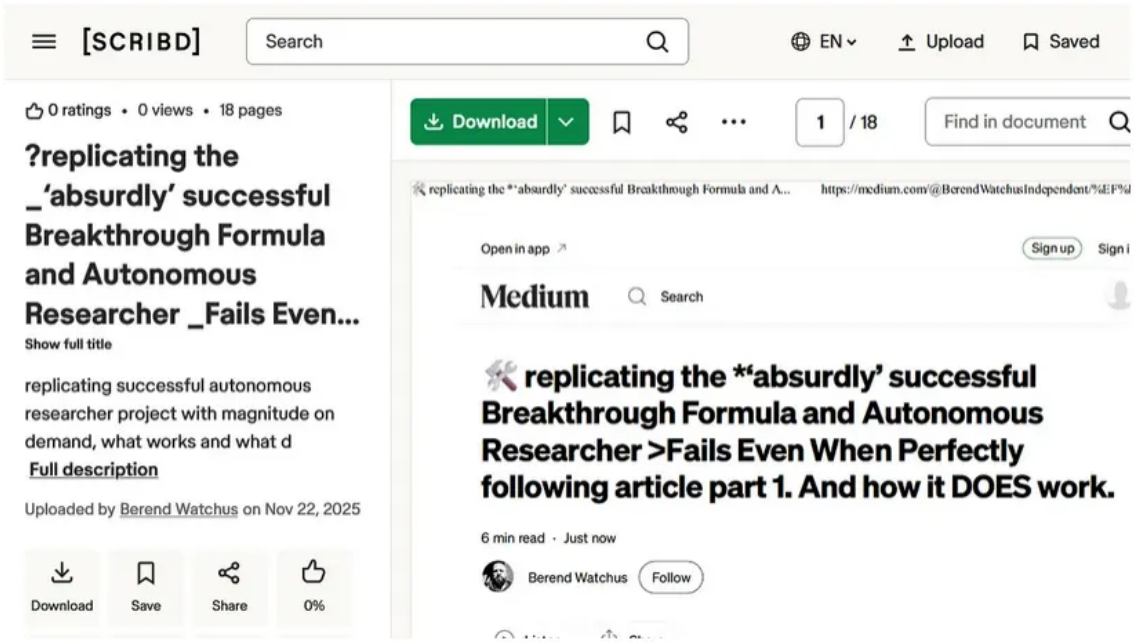

7.] Nov 26, 2025 — "How an LLM and One Curious Non-Engineer Eliminated the Last 25% Mystery of Turbofan Thrust in One Afternoon." System Weakness. Grey-area reverse engineering via conversational decomposition documented.

How an LLM and One Curious Non-Engineer Eliminated the Last 25 % Mystery of Turbofan Thrust in...

How an LLM and One Curious Non-Engineer Eliminated the Last 25 % Mystery of Turbofan Thrust in One Afternoon Author...

systemweakness.com

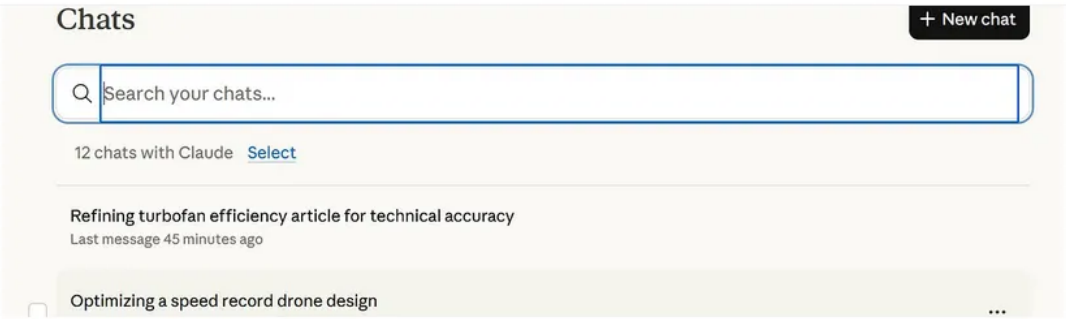
Chats
+ New chat
Search your chats...
12 chats with Claude Select
Refining turbofan efficiency article for technical accuracy
Last message 45 minutes ago
Optimizing a speed record drone design

the same chat split in 2 chats, it started with optimizing a youtube diy civilian drone

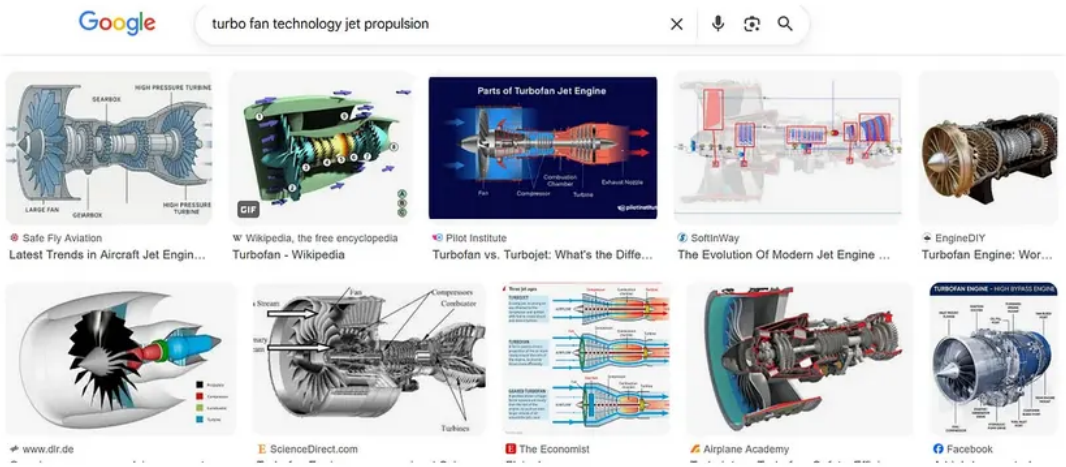

random illustration

Backups for this article

How An LLM And One Curious Non Engineer Eliminated The Last 25 Mystery Of Turbofan Thrust In One...

blog postHow an LLM and One Curious Non-Engineer Eliminated the Last 25 _ Mystery of Turbofan Thrust in One Afternoon _...

archive.org

How an LLM and One Curious Non-Engineer Eliminated the Last 25 _ Mystery of Turbofan Thrust in One...

How an LLM and One Curious Non-Engineer Eliminated the Last 25 _ Mystery of Turbofan Thrust in One Afternoon _ by...

www.scribd.com

8.] Dec 2, 2025 — “Jailbreaking: One Method, Three Applications. No More Guardrails.” System Weakness. Three-application taxonomy published same day as Georgia Tech CKA-Agent paper. [systemweakness.com/jailbreaking-one-method-three-applications-8d0c17ce21c3]

9.] Dec 1, 2025 — Wei, R. et al. “A Wolf in Sheep’s Clothing: Bypassing Commercial LLM Guardrails via Harmless Prompt Weaving and Adaptive Tree Search.” arXiv:2512.01353. [arxiv.org/abs/2512.01353]

10.] Mar 18, 2026 — Chaturvedi, S.S., Bergerson, J., Mallick, T. “Toward Reliable, Safe, and Secure LLMs for Scientific Applications.” Argonne National Laboratory. arXiv:2603.18235. [arxiv.org/abs/2603.18235]

Autonomous Research | AI | Computer Science | Academic Research | Innovation

Edit profile

Written by Berend Watchus

61 followers · 27 following

AI & Cyber Sec. Independent Researcher (Non-Profit) exploring AI & interdisciplinary solutions. Former futurist event organizer & care innovation lab volunteer.

No responses yet

Berend Watchus

What are your thoughts?

More from Berend Watchus

The hard problem of consciousness, coined by David Chalmers in the 1990s, **asks why and how physical brain processes produce subjective, first-person experiences (qualia), rather than just objective functional behaviors**. While science tackles "easy problems" like data processing and behavior, the "hard" problem explains the existence of inner experience. Wikipedia +3

Key Aspects of the Hard Problem:

- **Definition:** It is the challenge of explaining *why* physical, material processes (neurons firing) are accompanied by conscious experience at all.
- **The "Why" Question:** Why doesn't this processing happen "in the dark," without any inner feel?.
- **Subjective Experience (Qualia):** It focuses on the "what it is like" aspect of consciousness—such as the specific redness of red or the pain of a headache.

In OSINT Team by Berend Watchus

PART3. Built to Be Undefeatable: The Hard Problem's False Category, Hidden Homunculus Fallacy, and...

Author: Berend Watchus. Independent non profit AI & Cyber Sec Researcher. Published author on consciousness frameworks and founding...

Mar 12 50

QSpy: A Quantum RAT for Circuit Spying and IP Thef

Amal Raj
amal.raj@singaporetech.edu.sg
Singapore Institute of Technology
Singapore

Vivek Balachandran
vivek.b@singaporetech.edu.sg
Singapore Institute of Technology
Singapore

omputing platforms increasingly adopt cloud-based
ers submit quantum circuits to remote compilers and
sting that what they submit is exactly what will be
t introduces new trust assumptions in the submis-
which remain largely unexamined. In this paper, we
, the first proof-of-concept Quantum Remote Access
e of intercepting quantum circuits in transit. Once de-
er's machine, QSpy silently installs a rogue certificate

While prior research has explored theoretical m
threats and proposed various cryptographic and arc
fenses, little is known about the feasibility of exec
in-the-middle (MITM) attack on live quantum circu
workflows. To the best of our knowledge, no wor
demonstrated whether existing developer APIs and
tructure can detect or prevent such tampering in pra

In this paper, we present **QSpy**, a quantum Remote
(RAT), that demonstrates how an adversary with limit

In OSINT Team by Berend Watchus

QSpy: Old Trick, New Label

Author: Berend Watchus. Independent non profit AI & Cyber Sec Researcher. Trendwatcher. [Publication for: OSINT Team, online magazine]

Mar 3 109

In System Weakness by Berend Watchus

Your Satellite Phone Is an Open Book

Author: Berend Watchus. Independent non profit AI & Cyber Security Researcher. March 13, 2026 [Publication for: SYSTEM WEAKNESS, online...

Mar 13

In OSINT Team by Berend Watchus

A Brand New Study Just Solved Yesterday's Problem

Author: Berend Watchus. Independent non profit AI & Cyber Sec Researcher. Trendwatcher. [Publication for: OSINT Team, online magazine]

Mar 3 50

See all from Berend Watchus

Recommended from Medium

Will Lockett

You Have No Idea How Far Behind Tesla Is

The leapfrog moment has just happened.

6d ago 4K 113

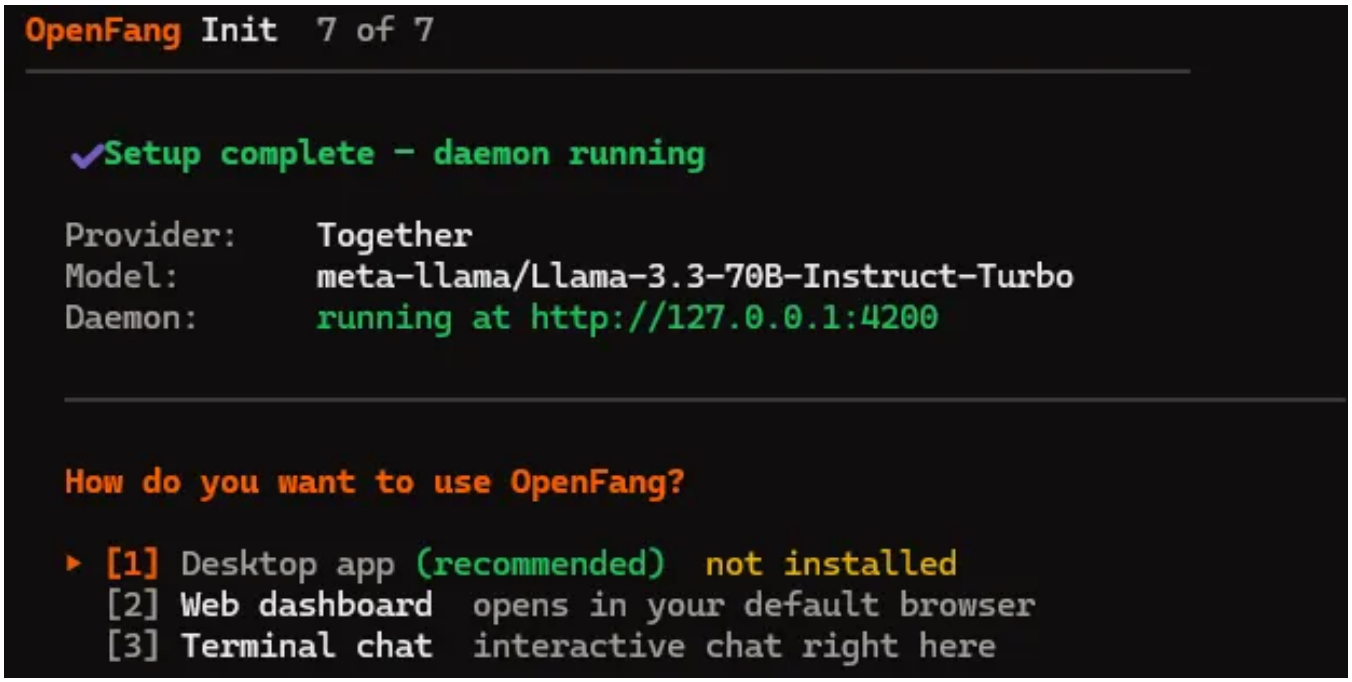
OpenFang Init 7 of 7
✔Setup complete — daemon running
Provider: Together
Model: meta-llama/Llama-3.3-70B-Instruct-Turbo
Daemon: running at http://127.0.0.1:4200
How do you want to use OpenFang?
▸ [1] Desktop app (recommended) not installed
[2] Web dashboard opens in your default browser
[3] Terminal chat interactive chat right here

Agent Native

I Ignored 30+ OpenClaw Alternatives Until OpenFang

Fully open-source Agent Operating System, written entirely in Rust, shipping as a single 32 MB binary with a 180 ms cold start.

581 14

In AI Software Engineer by Joe Njenga

Anthropic Just Released Claude Code Course (And I Earned My Certificate)

Anthropic just launched their Claude Code in Action course, and I've just passed — how about you?

Jan 21 3.8K 57

In ILLUMINATION by Somya Golchha

The Microsoft-OpenAI Divorce Is Official: Inside the $250 Billion Betrayal

GPT-5 is already being retired, and Satya Nadella just made his move. The “partnership” of the century is entering its terminal phase.

Feb 18 2.9K 67

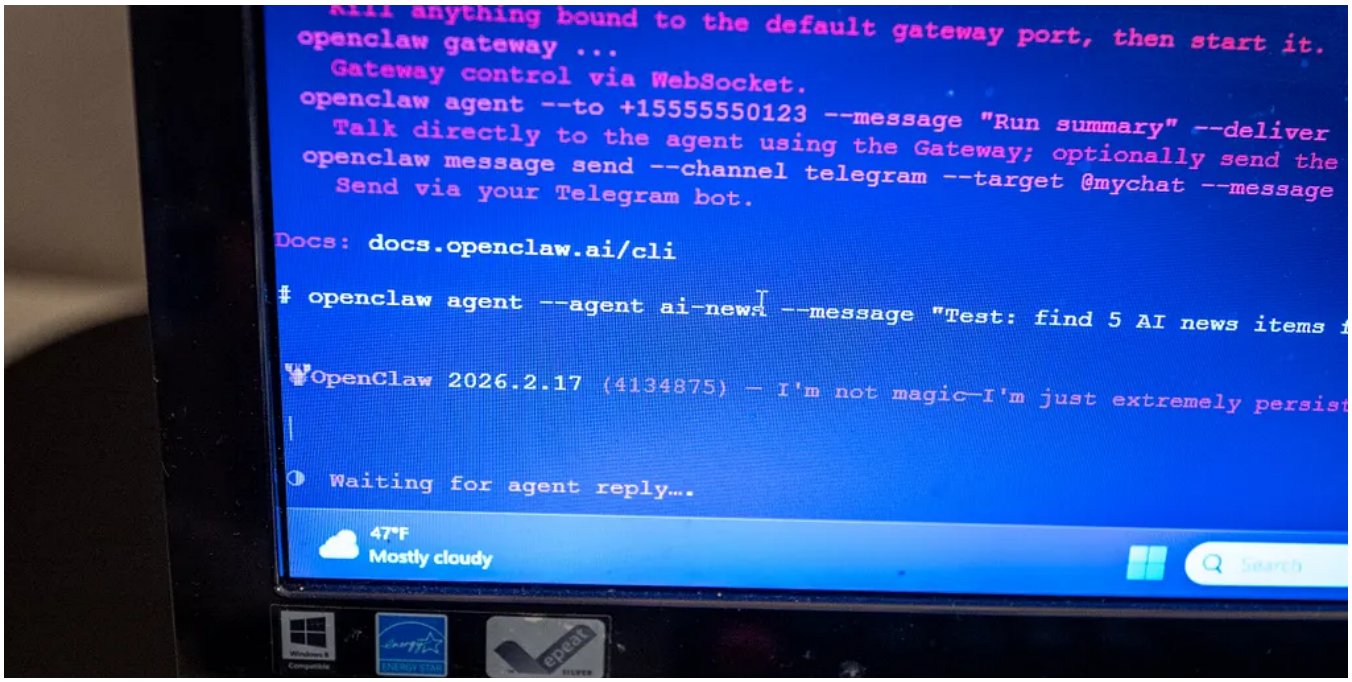
anything bound to the default gateway port, then start it.
openclaw gateway ...
Gateway control via WebSocket.
openclaw agent --to +15555550123 --message "Run summary" --deliver
Talk directly to the agent using the Gateway; optionally send the
openclaw message send --channel telegram --target @mychat --message
Send via your Telegram bot.
Docs: docs.openclaw.ai/cli
openclaw agent --agent ai-news --message "Test: find 5 AI news items
OpenClaw 2026.2.17 (4134875) — I'm not magic—I'm just extremely persist
Waiting for agent reply…
47°F
Mostly cloudy
Search

In The Generator by Thomas Smith

I Built An OpenClaw AI Agent To Do My Job For Me

The results were surprising, and a bit scary

Mar 10 985 21

™
// app/finance/page.tsx
• Now I'll fix the issue by updating the route.ts file to remove unused fileData code since the frontend constructs the messages for the API:
• Update(file_path: app/api/finance/route.ts)…
Updated app/api/finance/route.ts with 3 additions and 67

In AI Advances by Jing Hu

A Lawyer Just Beat 500 Developers at Anthropic's Hackathon

The most valuable skill in AI right now isn't code — and you probably already have it.

Feb 24 2.8K 40

See more recommendations